IN THE SPECIFICATION

On pages 23-34, replace the paragraph beginning at page 23, line 16 as follows:

> More than 50 percent of the available MIPS can be allocated to a single thread, although this will result in a non-deterministic inter-instruction ~~delay the~~ delay, the time between successive instructions from the same thread would not be the same. For some applications this varying inter-instruction delay is not a disadvantage. For example, a thread could be scheduled in slots 1, 2, 3, 5, 6, 7, 9, … to achieve 75 percent of the available MIPS of the CPU. One type of NRT thread scheduling rotates through each thread. That is, the threads are scheduled in order, with one instruction executed from each active thread. This type of semi-flexible scheduling permits non-real-time threads to be scheduled in the empty slots in the schedule, e.g., the quanta labeled "*" in Figure 7d, and in slots where the scheduled hard real-time thread is not active, e.g., in place of thread B if thread B is not active, as described above. This type of scheduling is sometimes referred to as "round robin" scheduling.